

Beam loading compensation for Super B-factories

Dmitry Teytelman

Outline

- I. What is beam loading and why is it a problem in high current storage rings?
- Synchronous phase transients
 - Longitudinal coupled-bunch instabilities

Outline

I. What is beam loading and why is it a problem in high current storage rings?

- Synchronous phase transients
- Longitudinal coupled-bunch instabilities

II. Methods for reducing the harmful effects of the beam loading

- Cavity parameter optimization
- RF feedback
- Coupled-bunch instability feedback

Outline

I. What is beam loading and why is it a problem in high current storage rings?

- Synchronous phase transients
- Longitudinal coupled-bunch instabilities

II. Methods for reducing the harmful effects of the beam loading

- Cavity parameter optimization
- RF feedback
- Coupled-bunch instability feedback

III. RF parameter selection procedure for SuperPEP

Outline

I. What is beam loading and why is it a problem in high current storage rings?

- Synchronous phase transients
- Longitudinal coupled-bunch instabilities

II. Methods for reducing the harmful effects of the beam loading

- Cavity parameter optimization
- RF feedback
- Coupled-bunch instability feedback

III. RF parameter selection procedure for SuperPEP

IV. Summary

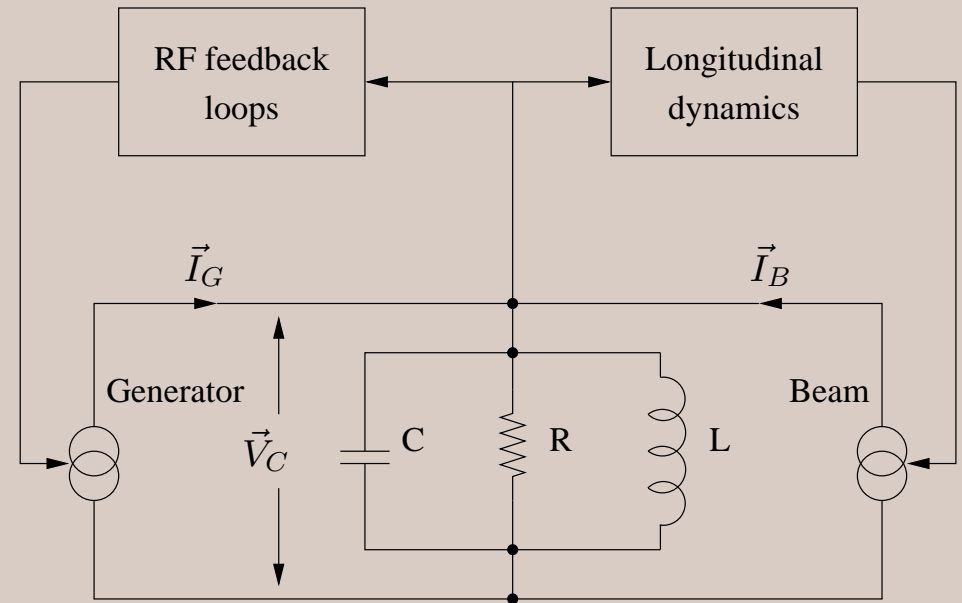
An RLC model of the RF cavity

We can model fundamental mode of the cavity as a parallel RLC circuit

$$Z(s) = \frac{2\sigma R s}{s^2 + 2\sigma s + \omega_r^2}$$

where $\sigma = \omega_r/(2Q)$ is the damping time of the cavity

The cavity is driven by two current sources: the generator (klystron) and the beam. Total cavity voltage is determined by the sum current and the cavity impedance.



When the beam current is small relative to the generator current - **light beam loading** - the cavity voltage is mostly defined by the generator current.

High beam current starts to affect strongly the cavity voltage thus creating a strong interaction between the RF system and the beam - **high beam loading condition**.

Think of the interaction as of a “feedback loop”: beam current source is affected by the cavity voltage, while that voltage depends on the beam current.

What does beam loading do?

Two most significant effects of high beam loading

- Synchronous phase transients due to uneven filling patterns
- Longitudinal coupled-bunch instabilities driven by the fundamental impedance


Synchronous phase transient: an example

A typical ring filling pattern includes a sizeable gap for the abort kicker and ion clearing

Gap is an amplitude modulation of the beam current

Translates into amplitude and phase modulations of the cavity voltage which induce a periodic synchronous phase transient.

PEP-II LER at 2.35 A, 3.8 MV




Effect of the synchronous phase transient

Different bunches see different RF voltage slopes and, therefore, have differing synchrotron tunes and bunch lengths.

Mismatched gap transients between the two rings shift the collision point position thus degrading the luminosity.

In the LFB front-end the transient appears as constant DC offsets of individual bunches. This has several consequences:



- Amplitude of the gap transient cannot exceed the full-scale peak-to-peak range of the phase detector used
- Largest expected gap transient amplitude sets the feedback front-end gain - need to properly detect motion for the bunches at the extremes of the transient.
- Phase detector gain rolls off as $\cos(M\phi)$ where M is the detection harmonic

What does beam loading do?

Two most significant effects of high beam loading

- ✓ Synchronous phase transients due to uneven filling patterns
- Longitudinal coupled-bunch instabilities driven by the fundamental impedance


Fundamental impedance and coupled-bunch instabilities

The growth rate of eigenmode -1 is proportional to the difference between the real parts of the impedance at $\omega_{\text{rf}} - \omega_{\text{rev}} + \omega_s$ and $\omega_{\text{rf}} + \omega_{\text{rev}} + \omega_s$ and

When the cavity is at resonance that difference is very small

However with increasing beam current the cavity center frequency is detuned below the RF frequency causing larger and larger asymmetries

When the detuning is comparable to the revolution frequency the instability growth rates become too fast to control




Fundamental impedance and coupled-bunch instabilities

The growth rate of eigenmode -1 is proportional to the difference between the real parts of the impedance at $\omega_{\text{rf}} - \omega_{\text{rev}} + \omega_s$ and $\omega_{\text{rf}} + \omega_{\text{rev}} + \omega_s$ and

When the cavity is at resonance that difference is very small

However with increasing beam current the cavity center frequency is detuned below the RF frequency causing larger and larger asymmetries

When the detuning is comparable to the revolution frequency the instability growth rates become too fast to control




Fundamental impedance and coupled-bunch instabilities

The growth rate of eigenmode -1 is proportional to the difference between the real parts of the impedance at $\omega_{\text{rf}} - \omega_{\text{rev}} + \omega_s$ and $\omega_{\text{rf}} + \omega_{\text{rev}} + \omega_s$ and

When the cavity is at resonance that difference is very small

However with increasing beam current the cavity center frequency is detuned below the RF frequency causing larger and larger asymmetries

When the detuning is comparable to the revolution frequency the instability growth rates become too fast to control




Fundamental impedance and coupled-bunch instabilities

The growth rate of eigenmode -1 is proportional to the difference between the real parts of the impedance at $\omega_{\text{rf}} - \omega_{\text{rev}} + \omega_s$ and $\omega_{\text{rf}} + \omega_{\text{rev}} + \omega_s$ and

When the cavity is at resonance that difference is very small

However with increasing beam current the cavity center frequency is detuned below the RF frequency causing larger and larger asymmetries

When the detuning is comparable to the revolution frequency the instability growth rates become too fast to control




Fundamental impedance and coupled-bunch instabilities

The growth rate of eigenmode -1 is proportional to the difference between the real parts of the impedance at $\omega_{\text{rf}} - \omega_{\text{rev}} + \omega_s$ and $\omega_{\text{rf}} + \omega_{\text{rev}} + \omega_s$ and

When the cavity is at resonance that difference is very small

However with increasing beam current the cavity center frequency is detuned below the RF frequency causing larger and larger asymmetries

When the detuning is comparable to the revolution frequency the instability growth rates become too fast to control



What does beam loading do?


Two most significant effects of high beam loading

- ✓ Synchronous phase transients due to uneven filling patterns
- ✓ Longitudinal coupled-bunch instabilities driven by the fundamental impedance
- An obvious question: why is this a problem? Couldn't one just build bigger feedback kickers, buy more power amplifiers and fix the instability problem with feedback?

Limits on achievable instability control

Controllable growth rates are limited by

- Maximum stable loop gain - depends on controller design, total loop delay.
- Maximum usable loop gain - gain that provides the largest damping. Depends on the same parameters as the maximum stable gain, but is significantly lower.
- Noise floor at the ADC - depends on RF-driven noise level, front-end electronics design
- Transient sensitivity - effect of injection and RF transients on longitudinal control. The sensitivity can be reduced by increasing kicker voltage.




For a conventional instability feedback the minimum group delay is one turn.

Experience with the low group-delay feedback channel in PEP-II (poster MPPP007) shows that a one-turn delay channel can reach 10 ms^{-1} damping at the 6 kHz synchrotron frequency.

Limits on achievable instability control

Controllable growth rates are limited by

- Maximum stable loop gain - depends on controller design, total loop delay.
- Maximum usable loop gain - gain that provides the largest damping. Depends on the same parameters as the maximum stable gain, but is significantly lower.
- Noise floor at the ADC - depends on RF-driven noise level, front-end electronics design
- Transient sensitivity - effect of injection and RF transients on longitudinal control. The sensitivity can be reduced by increasing kicker voltage.




For a conventional instability feedback the minimum group delay is one turn.

Experience with the low group-delay feedback channel in PEP-II (poster MPPP007) shows that a one-turn delay channel can reach 10 ms⁻¹ damping at the 6 kHz synchrotron frequency.

Limits on achievable instability control

Controllable growth rates are limited by

- Maximum stable loop gain - depends on controller design, total loop delay.
- Maximum usable loop gain - gain that provides the largest damping. Depends on the same parameters as the maximum stable gain, but is significantly lower.
- Noise floor at the ADC - depends on RF-driven noise level, front-end electronics design
- Transient sensitivity - effect of injection and RF transients on longitudinal control. The sensitivity can be reduced by increasing kicker voltage.




For a conventional instability feedback the minimum group delay is one turn.

Experience with the low group-delay feedback channel in PEP-II (poster MPPP007) shows that a one-turn delay channel can reach 10 ms^{-1} damping at the 6 kHz synchrotron frequency.

Limits on achievable instability control

Controllable growth rates are limited by

- Maximum stable loop gain - depends on controller design, total loop delay.
- Maximum usable loop gain - gain that provides the largest damping. Depends on the same parameters as the maximum stable gain, but is significantly lower.
- Noise floor at the ADC - depends on RF-driven noise level, front-end electronics design
- Transient sensitivity - effect of injection and RF transients on longitudinal control. The sensitivity can be reduced by increasing kicker voltage.




For a conventional instability feedback the minimum group delay is one turn.

Experience with the low group-delay feedback channel in PEP-II (poster MPPP007) shows that a one-turn delay channel can reach 10 ms^{-1} damping at the 6 kHz synchrotron frequency.

Limits on achievable instability control

Controllable growth rates are limited by

- Maximum stable loop gain - depends on controller design, total loop delay.
- Maximum usable loop gain - gain that provides the largest damping. Depends on the same parameters as the maximum stable gain, but is significantly lower.
- Noise floor at the ADC - depends on RF-driven noise level, front-end electronics design
- Transient sensitivity - effect of injection and RF transients on longitudinal control. The sensitivity can be reduced by increasing kicker voltage.



For a conventional instability feedback the minimum group delay is one turn.

Experience with the low group-delay feedback channel in PEP-II (poster MPPP007) shows that a one-turn delay channel can reach 10 ms⁻¹ damping at the 6 kHz synchrotron frequency.

What does beam loading do?

Two most significant effects of high beam loading

- ✓ Synchronous phase transients due to uneven filling patterns
- ✓ Longitudinal coupled-bunch instabilities driven by the fundamental impedance
- ✓ An obvious question: why is this a problem? Couldn't one just build bigger feedback kickers, buy more power amplifiers and fix the instability problem with feedback?

What can we do to reduce these harmful effects to manageable levels?


Compensating the phase transient

The vector diagram shows that the beam phase modulation has to be compensated by a significant amplitude modulation of the generator current

That modulation directly transforms into a significantly increased RF power requirement!

What else can one do to reduce the gap transient?

- Increase cavity stored energy
- Reduce the fill pattern gap
- PEP-II started with a 5% abort gap, upgraded to 2.5% gap, expect to go to 1.5%
- An interesting idea to explore is how much of a gap transient can be compensated by only **phase** modulating the RF
- A calculation made by P. Wilson in 1992 shows that a 12 degree transient in PEP-II HER could be reduced to 3 degrees with phase modulation only.



What does beam loading do?

Two most significant effects of high beam loading

- ✓ Synchronous phase transients due to uneven filling patterns
- ✓ Longitudinal coupled-bunch instabilities driven by the fundamental impedance
- ✓ An obvious question: why is this a problem? Couldn't one just build bigger feedback kickers, buy more power amplifiers and fix the instability problem with feedback?

What can we do to reduce these harmful effects to manageable levels?

- ✓ Synchronous phase transients
 - ✓ Increase cavity stored energy, reduce the gap length, possibly use RF phase modulation
- Longitudinal instabilities due to the fundamental impedance of the RF cavity

Longitudinal coupled-bunch instabilities

Methods for fighting the longitudinal coupled-bunch instabilities due to the fundamental impedance of the RF cavity include

- Reduce the cavity detuning by increasing stored energy
- Minimize the number of RF cavities to reduce the total impedance
- RF feedback to reduce the cavity impedance seen by the beam
- Instability feedback to deal with the residual growth rates

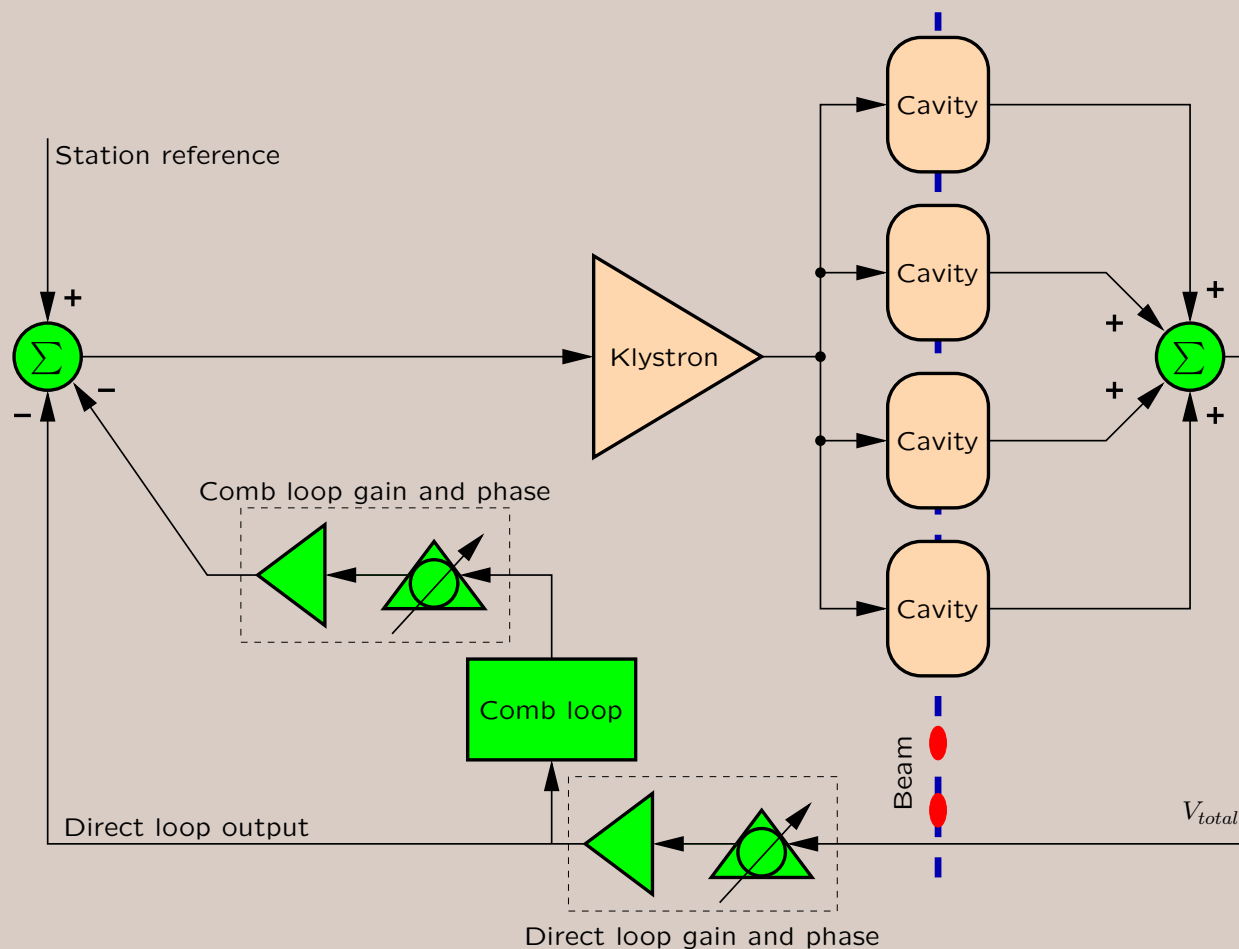
There is no single “magic” solution - a successful design must include all of the above to suppress the instabilities

Example: SuperKEKB

- High stored energy in superconducting and ARES cavities
- ARES cavity upgrade for increased stored energy
- Mode -1 LLRF feedback
- Longitudinal bunch-by-bunch feedback to control the residual growth rates

PEP-II fast impedance control loops: topology

The most important elements of the impedance controlling feedback loops are shown. The direct feedback loop uses the cavity vector sum signal (a complex signal), scaled in magnitude and rotated in phase as an input to a reference summing node. The comb loop (a periodic IIR filter) uses the direct loop output via the comb filter, scaled and rotated, as a summing input.



The overall action of this feedback topology is to keep the combined direct and comb outputs exactly equal to the station reference - any error signal is amplified via the klystron and cavity path. The overall station cavity magnitude and phase are set via this reference.

PEP-II low-level RF feedback: impedances and growth rates

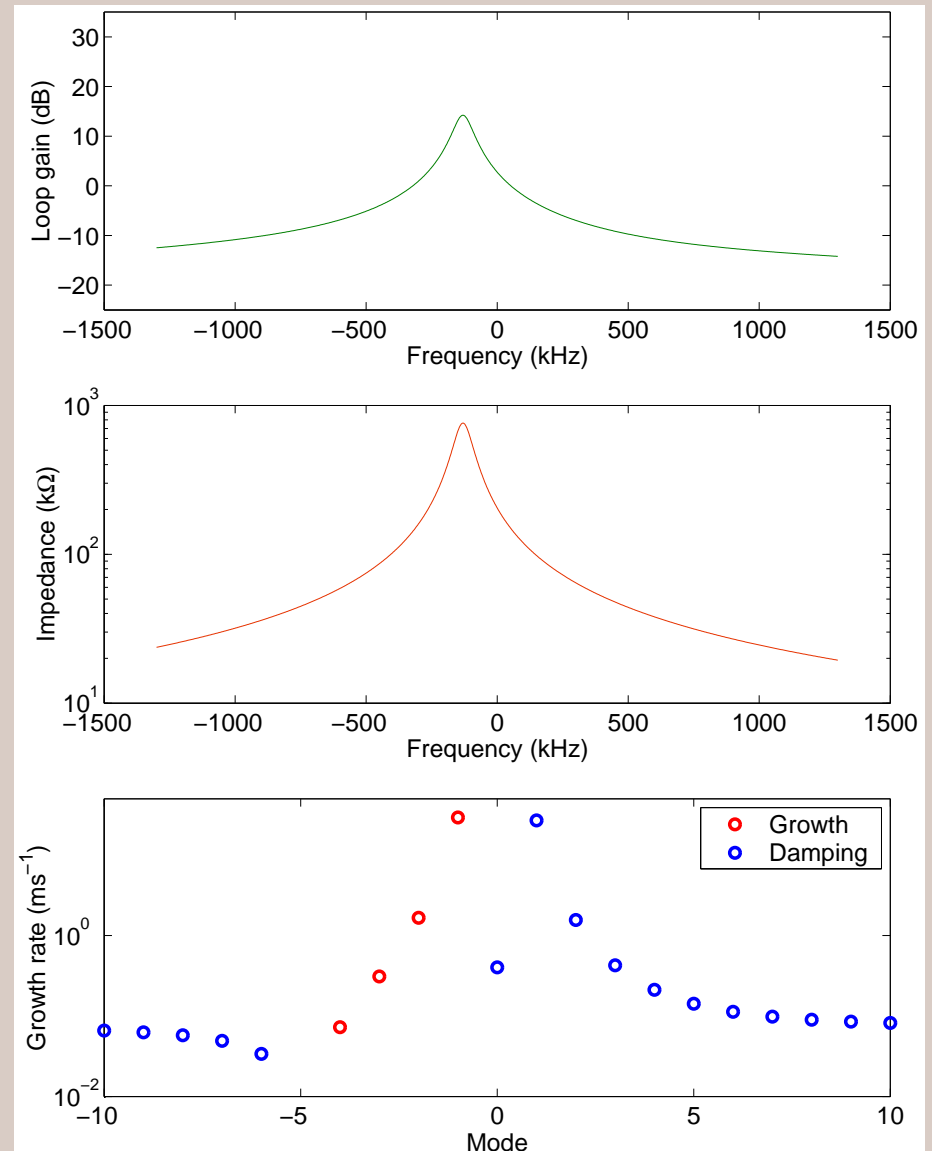
Two feedback loops are used in PEP-II to reduce the fundamental impedance acting on the beam: direct and comb.

Direct loop is a proportional feedback loop around the cavity. Closing the direct feedback loop reduces the effective impedance seen by the beam and lowers the growth rates.

To reduce the growth rates further we add the comb filter with narrow gain peaks at synchrotron sidebands.

Expected growth rates shown here are computed using a linear transfer function model of the RF feedback system.

According to the linear model the growth rate reduction is two orders of magnitude, from 30 to 0.35 ms^{-1}



PEP-II low-level RF feedback: impedances and growth rates

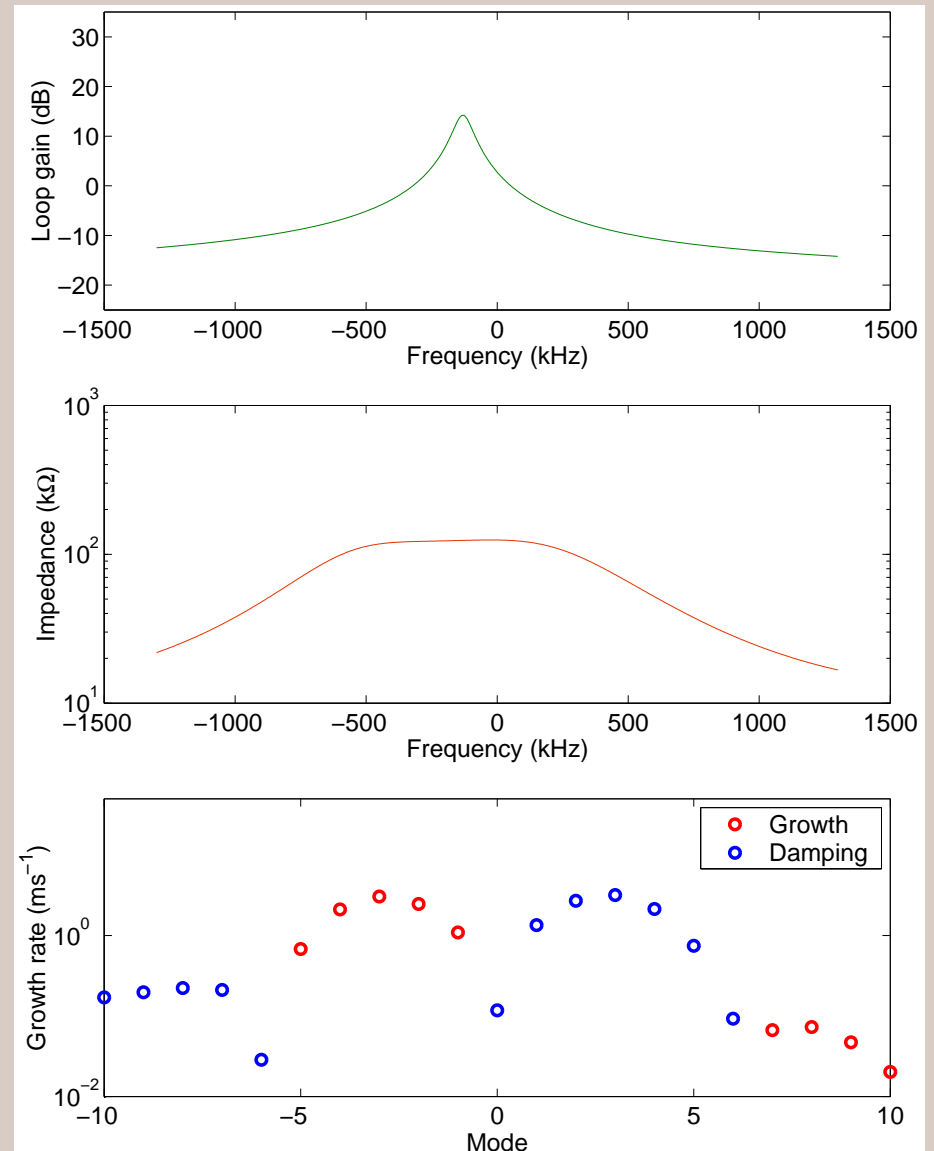
Two feedback loops are used in PEP-II to reduce the fundamental impedance acting on the beam: direct and comb.

Direct loop is a proportional feedback loop around the cavity. Closing the direct feedback loop reduces the effective impedance seen by the beam and lowers the growth rates.

To reduce the growth rates further we add the comb filter with narrow gain peaks at synchrotron sidebands.

Expected growth rates shown here are computed using a linear transfer function model of the RF feedback system.

According to the linear model the growth rate reduction is two orders of magnitude, from 30 to 0.35 ms^{-1}



PEP-II low-level RF feedback: impedances and growth rates

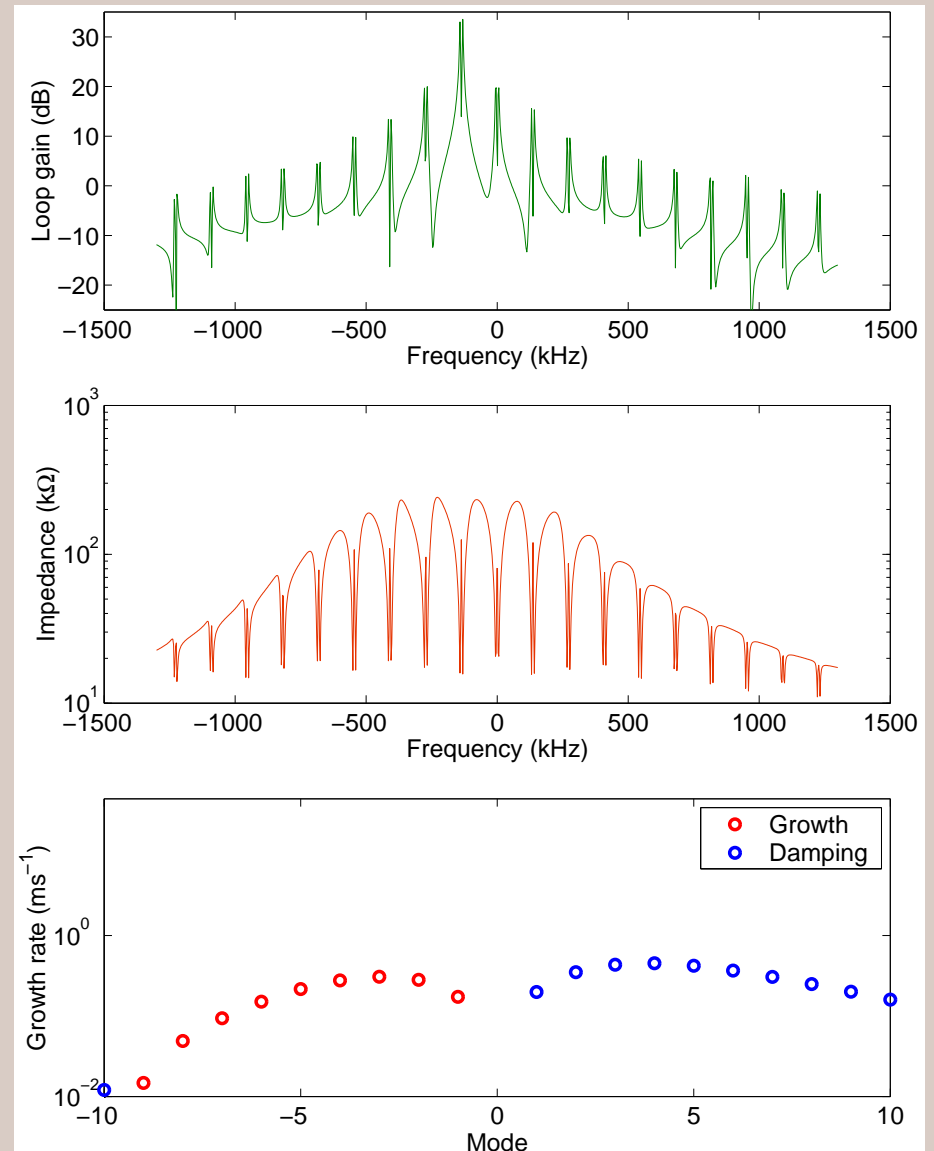
Two feedback loops are used in PEP-II to reduce the fundamental impedance acting on the beam: direct and comb.

Direct loop is a proportional feedback loop around the cavity. Closing the direct feedback loop reduces the effective impedance seen by the beam and lowers the growth rates.

To reduce the growth rates further we add the comb filter with narrow gain peaks at synchrotron sidebands.

Expected growth rates shown here are computed using a linear transfer function model of the RF feedback system.

According to the linear model the growth rate reduction is two orders of magnitude, from 30 to 0.35 ms^{-1}



What does beam loading do?

Two most significant effects of high beam loading

- ✓ Synchronous phase transients due to uneven filling patterns
- ✓ Longitudinal coupled-bunch instabilities driven by the fundamental impedance
- ✓ An obvious question: why is this a problem? Couldn't one just build bigger feedback kickers, buy more power amplifiers and fix the instability problem with feedback?

What can we do to reduce these harmful effects to manageable levels?

- ✓ Synchronous phase transients
 - ✓ Increase cavity stored energy, reduce the gap length, possibly use RF phase modulation
- ✓ Longitudinal instabilities due to the fundamental impedance of the RF cavity
 - ✓ Increase cavity stored energy, minimize the number of cavities, LLRF feedback, coupled-bunch instability feedback

Minimizing the fundamental impedance

Minimize the number of cavities

Keep the detuning low

$$\omega_D = \left| \frac{\omega_{\text{rf}} I_0 R}{V_c Q} \cos \phi_B \right| \approx \frac{\omega_{\text{rf}} I_0 R}{V_c Q}$$

To achieve low detuning

- Need low R/Q
- It is desirable to operate the cavities at as high a voltage as possible

In cavity design low R/Q leads to lower achievable cavity voltage.

Might be useful to minimize the quantity $\frac{1}{V_c} \frac{R}{Q}$

Based on these requirements we come up with a “cookbook” procedure for selecting ring RF parameters.

Determining SuperPEP ring parameters

Start from the achievable cavity parameters:

- Power coupled to each cavity P_g
- Maximum cavity voltage V_c

Compute the total beam power requirements due to the synchrotron radiation, resistive wall and HOM losses.

Minimum number of cavities N_c is determined by the ratio of the beam power to the power delivered to the beam per cavity

Set the total RF voltage to the largest achievable value $N_c V_c$

$$\text{From } \sigma_z = \frac{\alpha \delta_E c}{\omega_s} \text{ and } \omega_s^2 = \frac{\alpha e \omega_{\text{rf}}}{E_0 T_0} V_G \text{ we get } \alpha = \frac{\omega_{\text{rf}} e V_G \sigma_z^2}{E_0 T_0 \delta_E^2 c^2}$$

Desired bunch length and gap voltage set the momentum compaction for the ring. For constant bunch length the ratio α/V_G is fixed. If we push the cavity voltage higher the momentum compaction has to increase as well leading to a linear increase in the synchrotron frequency.

Determining SuperPEP ring parameters: assumptions

Only superconducting cavities are considered

- Conventional normal conducting cavities are unfeasible - very large wall and HOM losses, huge detuning frequencies
- Energy storage cavities have several disadvantages relative to the superconducting cavities
 - Wall power loss - at the same generator power one will need more energy storage cavities than superconducting ones
 - Relatively low cavity voltage - requires matching low momentum compaction which might be difficult to achieve

Synchronous phase angle is very close to π - quite reasonable for the large overvoltage factors being considered

We can couple 1 MW into each cavity

Maximum cavity voltage is 1.25 MV

- A reasonable assumption for the cavities with R/Q of 5Ω , might be too conservative for higher R/Q .

Parameter decision procedure example

LER at 3.5 GeV and 15.5 A

Synchrotron radiation loss of 15.04 MW

Resistive wall loss of 2.76 MW

HOM loss (excluding RF cavities) of 2.32 MW: total of 20.12 MW

Power delivered to the beam per cavity (loss factor of 0.36 V/pC) is 908 kW

Need 22 cavities

At 1.25 MV per cavity total gap voltage is 27.5 MV

Assuming fractional energy spread $\delta_E = 8 \cdot 10^{-4}$ for $\sigma_z = 1.8$ mm we get

$$\alpha = 3.6 \cdot 10^{-4}$$

$$f_s = 7.65 \text{ kHz}$$

Cavity options under consideration

Cavity	$R/Q, \Omega$	I_0, A	$\alpha, 10^{-4}$	N_c	$\Delta f, \text{kHz}$	$P_{\text{HOM}}, \text{kW}$	P_b, kW
SC952	30	15.5	3.6	23	353.6	92	908
SC952a	12		3.6	23	141.7	79	921
SC952b	5		3.6	22	60.7	72	928
SC952	30	23	6.9	42	524.7	202	798
SC952a	12		6.7	41	210.2	174	826
SC952b	5		6.6	40	90.1	158	842

Cavity options under consideration

Cavity	$R/Q, \Omega$	I_0, A	$\alpha, 10^{-4}$	N_c	$\Delta f, \text{kHz}$	$P_{\text{HOM}}, \text{kW}$	P_b, kW
SC952	30	15.5	3.6	23	353.6	92	908
SC952a	12		3.6	23	141.7	79	921
SC952b	5		3.6	22	60.7	72	928
SC952	30	23	6.9	42	524.7	202	798
SC952a	12		6.7	41	210.2	174	826
SC952b	5		6.6	40	90.1	158	842

For high R/Q the detuning is very large - from 2.5 to almost 4 revolution harmonics

Cavity options under consideration

Cavity	$R/Q, \Omega$	I_0, A	$\alpha, 10^{-4}$	N_c	$\Delta f, \text{kHz}$	$P_{\text{HOM}}, \text{kW}$	P_b, kW
SC952	30	15.5	3.6	23	353.6	92	908
SC952a	12		3.6	23	141.7	79	921
SC952b	5		3.6	22	60.7	72	928
SC952	30	23	6.9	42	524.7	202	798
SC952a	12		6.7	41	210.2	174	826
SC952b	5		6.6	40	90.1	158	842

For high R/Q the detuning is very large - from 2.5 to almost 4 revolution harmonics

At the other end of the spectrum low R/Q leads to detuning frequencies under one revolution harmonic

Cavity options under consideration

Cavity	$R/Q, \Omega$	I_0, A	$\alpha, 10^{-4}$	N_c	$\Delta f, \text{kHz}$	$P_{\text{HOM}}, \text{kW}$	P_b, kW
SC952	30	15.5	3.6	23	353.6	92	908
SC952a	12		3.6	23	141.7	79	921
SC952b	5		3.6	22	60.7	72	928
SC952	30	23	6.9	42	524.7	202	798
SC952a	12		6.7	41	210.2	174	826
SC952b	5		6.6	40	90.1	158	842

For high R/Q the detuning is very large - from 2.5 to almost 4 revolution harmonics

At the other end of the spectrum low R/Q leads to detuning frequencies under one revolution harmonic

HOM power loss ranges from 7% to 20% of the input power as a function of the loss factor and the beam current.

Growth rates for different cavity designs

Here we consider three RF system configurations

- PEP-II-like LLRF feedback (direct loop + comb filter)
- Same plus klystron linearizer for better impedance reduction
- No RF feedback for cavity SC952b (R/Q of 5)

Cavity	I_0 , A	Δf , kHz	R_{tot} , k Ω	Mode	Rate (sat), ms ⁻¹	Rate (lin), ms ⁻¹
SC952	15.5	353.6	1563	-3	10.58	2.12
SC952a		141.7	584	-3	3.95	0.79
SC952b		60.7	31.7	-1	0.43	
SC952	23	524.7	2986	-2	30	6
SC952a		210.2	1200	-3	12.05	2.41
SC952b		90.1	284	-1	5.7	

From the operational experience in many storage rings we believe that rates under 5 ms⁻¹ should be controllable, higher growth rates start eroding the stability margin

Cavity design comparison

Cavity	I_0 , A	Δf , kHz	R_{tot} , k Ω	Mode	Rate (sat), ms ⁻¹	Rate (lin), ms ⁻¹
SC952	15.5	353.6	1563	-3	10.58	2.12
SC952a		141.7	584	-3	3.95	0.79
SC952b		60.7	31.7	-1	0.43	
SC952	23	524.7	2986	-2	30	6
SC952a		210.2	1200	-3	12.05	2.41
SC952b		90.1	284	-1	5.7	

The R/Q of 30 Ω only works if we have linearized klystrons. Even then it is just marginal at 10^{36}

Cavity design comparison

Cavity	I_0 , A	Δf , kHz	R_{tot} , k Ω	Mode	Rate (sat), ms ⁻¹	Rate (lin), ms ⁻¹
SC952	15.5	353.6	1563	-3	10.58	2.12
SC952a		141.7	584	-3	3.95	0.79
SC952b		60.7	31.7	-1	0.43	
SC952	23	524.7	2986	-2	30	6
SC952a		210.2	1200	-3	12.05	2.41
SC952b		90.1	284	-1	5.7	

For the R/Q of 12 Ω existing LLRF feedback structure would be sufficient at 15.5 A, but at 23 A we would need to linearize the klystrons.

Currently a preferred choice as a good compromise between fundamental-driven growth rates and the aggressiveness in lowering R/Q .

Cavity design comparison

Cavity	I_0 , A	Δf , kHz	R_{tot} , k Ω	Mode	Rate (sat), ms ⁻¹	Rate (lin), ms ⁻¹
SC952	15.5	353.6	1563	-3	10.58	2.12
SC952a		141.7	584	-3	3.95	0.79
SC952b		60.7	31.7	-1	0.43	
SC952	23	524.7	2986	-2	30	6
SC952a		210.2	1200	-3	12.05	2.41
SC952b		90.1	284	-1	5.7	

Since this cavity design was evaluated without feedback there are several unique advantages to that approach

- LLRF feedback system is eliminated.
- Klystrons can be fully saturated leading to better power efficiency.

Growth rate is relatively high at 23 A - marginal control.

- Adding LLRF feedback drops the growth rate to 3.48 ms⁻¹ (0.7 ms⁻¹)

Summary

High current storage rings must pay careful attention to the harmful beam loading effects

Longitudinal coupled-bunch instabilities due to the cavity fundamental impedance to large extent define the RF system design for a highly beam loaded storage ring

Reducing the growth rates of such instabilities to a manageable level will most likely involve a combination of several methods

- Impedance minimization techniques
- Number of cavities
- Cavity detuning (stored energy)
- LLRF feedback

Superconducting cavities are the optimal choice for minimizing the instability driving impedance.

Cavity stored energy increase of 4 to 10 times relative to the existing superconducting cavities is required to produce acceptable growth rates at the high beam currents proposed for SuperPEP.

Acknowledgments

Many thanks to R. Akre, M. Browne, J. Byrd, P. Corredoura, A. Drago, J. Fox, A. Gallo, H. Hindi, S. Prabhakar, F. Pedersen, H. Schwarz, M. Serio, R. Tighe, and D. Van Winkle for numerous valuable discussions and ideas.

Special thanks to P. McIntosh, A. Novokhatski, J. Seeman, M. Sullivan, and U. Wienands for discussions of SuperPEP RF and accelerator parameters.

I would also like to thank the Accelerator Department at SLAC for an opportunity to think about SuperPEP RF design choices.

I am also very thankful to Boris Podobedov and Dan Van Winkle for suffering through a preliminary version of this talk and for their suggestions.

Work supported by U.S. Department of Energy contract DE-AC02-76SF00515

